

AMENDMENTS TO THE CLAIMS

Please amend the claims as follows.

1. – 25. (Canceled)

26. (Currently Amended) A method performed by an operating system executing within a computer system having different types of physical processing modules, the method comprising:

receiving a first set of parametric information pertaining to a first physical processing module (PPM), the first PPM having a first architecture;

receiving a second set of parametric information pertaining to a second PPM, the second PPM having a second architecture, wherein the first architecture is second PPM has a different from the second architecture than the first PPM;

constructing a first abstraction of the first PPM based, at least partially, upon the first set of parametric information, the first abstraction comprising an indication of how many logical processing entities are provided by the first PPM, the first abstraction further comprising operational information indicating one or more operational characteristics of the first PPM; ~~[[and]]~~

constructing a second abstraction of the second PPM based, at least partially, upon the second set of parametric information, the second abstraction comprising an indication of how many logical processing entities are provided by the second PPM, the second abstraction further comprising operational information indicating one or more operational characteristics of the second PPM; and

determining, based on at least one selected from a group consisting of the operational information in the first abstraction and the operational information in the second abstraction, whether to dispatch an execution thread for execution to one selected from a group consisting of one of the logical processing entities of the first PPM and one of the logical processing entities of the second PPM;

wherein the second abstraction is different from the first abstraction to reflect the differences between the first architecture of the first PPM and the second different architecture of the second PPM.

27. (Original) The method of claim 26, wherein the first PPM comprises a first physical processing core which supports a single hardware thread, wherein the second PPM comprises a second physical processing core which supports a plurality of hardware threads, wherein constructing the first abstraction comprises indicating the first physical processing core as a single logical processing entity, and wherein constructing the second abstraction comprises indicating the second physical processing core as a plurality of logical processing entities.
28. (Original) The method of claim 27, wherein constructing the second abstraction comprises indicating that the plurality of logical processing entities corresponding to the second physical processing core share one or more resources of the second PPM.
29. – 53. (Canceled)
54. (Currently Amended) A computer readable storage medium carrying instructions for an operating system, wherein the instructions, when executed in a computer system having different types of physical processing modules, cause the computer system to perform the operations of:
- receiving a first set of parametric information pertaining to a first physical processing module (PPM), the first PPM having a first architecture;
 - receiving a second set of parametric information pertaining to a second PPM, the second PPM having a second architecture, wherein the first architecture is second-PPM-has-a different from the second architecture ~~than the first PPM~~;
 - constructing a first abstraction of the first PPM based, at least partially, upon the first set of parametric information, the first abstraction comprising an indication of how many logical processing entities are provided by the first PPM, the first abstraction further comprising operational information indicating one or more operational characteristics of the first PPM; [[and]]

constructing a second abstraction of the second PPM based, at least partially, upon the second set of parametric information, the second abstraction comprising an indication of how many logical processing entities are provided by the second PPM, the second abstraction further comprising operational information indicating one or more operational characteristics of the second PPM; and

determining, based on at least one selected from a group consisting of the operational information in the first abstraction and the operational information in the second abstraction, whether to dispatch an execution thread for execution to one selected from a group consisting of one of the logical processing entities of the first PPM and one of the logical processing entities of the second PPM;

wherein the second abstraction is different from the first abstraction to reflect the differences between the first architecture of the first PPM and the second ~~different~~ architecture of the second PPM.

55. (Previously Presented) The computer readable storage medium of claim 54, wherein the first PPM comprises a first physical processing core which supports a single hardware thread, wherein the second PPM comprises a second physical processing core which supports a plurality of hardware threads, wherein constructing the first abstraction comprises indicating the first physical processing core as a single logical processing entity, and wherein constructing the second abstraction comprises indicating the second physical processing core as a plurality of logical processing entities.

56. (Previously Presented) The computer readable storage medium of claim 55, wherein constructing the second abstraction comprises indicating that the plurality of logical processing entities corresponding to the second physical processing core share one or more resources of the second PPM.

57. – 81. (Canceled)

82. (Currently Amended) A computer system comprising:

a first physical processing module (PPM) having a first architecture;

a second PPM having a second architecture, wherein the first architecture is second PPM has
a different from the second architecture than the first PPM; and

an operating system executing within the computer system, wherein the operating system performs the operations of:

receiving a first set of parametric information pertaining to the first PPM;

receiving a second set of parametric information pertaining to the second PPM,;

constructing a first abstraction of the first PPM based, at least partially, upon the first set of parametric information, the first abstraction comprising an indication of how many logical processing entities are provided by the first PPM, the first abstraction further comprising operational information indicating one or more operational characteristics of the first PPM; [[and]]

constructing a second abstraction of the second PPM based, at least partially, upon the second set of parametric information, the second abstraction comprising an indication of how many logical processing entities are provided by the second PPM, the second abstraction further comprising operational information indicating one or more operational characteristics of the second PPM; and

determining, based on at least one selected from a group consisting of the operational information in the first abstraction and the operational information in the second abstraction, whether to dispatch an execution thread for execution to one selected from a group consisting of one of the logical processing entities of the first PPM and one of the logical processing entities of the second PPM;

wherein the second abstraction is different from the first abstraction to reflect the differences between the first architecture of the first PPM and the second different architecture of the second PPM.

83. (Previously Presented) The computer system of claim 82, wherein the first PPM comprises a first physical processing core which supports a single hardware thread, wherein the second PPM comprises a second physical processing core which supports a plurality of hardware threads, wherein constructing the first abstraction comprises indicating the first physical processing core as a single logical processing entity, and wherein constructing the second abstraction comprises indicating the second physical processing core as a plurality of logical processing entities.
84. (Previously Presented) The computer system of claim 83, wherein constructing the second abstraction comprises indicating that the plurality of logical processing entities corresponding to the second physical processing core share one or more resources of the second PPM.
85. (Previously Presented) The method of claim 26,
 wherein the first PPM has an n number of processing cores, where n is an integer greater than one, each of the n processing cores capable of supporting one hardware thread;
 and
 wherein the second PPM has a single processing core capable of supporting an m number of hardware threads, where m is an integer greater than one.
86. (Previously Presented) The method of claim 85, wherein the first abstraction indicates that the first PPM provides an n number of logical processing entities, and wherein the second abstraction indicates that the second PPM provides an m number of logical processing entities.
87. (Previously Presented) The method of claim 86, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the n logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by the m logical processing entities of the second PPM.
88. (Previously Presented) The method of claim 87, wherein the resource sharing information of the first abstraction indicates that the n logical processing entities of the first PPM share a cache, and wherein the resource sharing information of the second abstraction indicates that the m

logical processing entities of the second PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.

89. (Previously Presented) The method of claim 26,

wherein the first PPM has an n number of processing cores, where n is an integer greater than one, each of the n processing cores capable of supporting one hardware thread; and

wherein the second PPM has an m number of processing cores, wherein m is an integer greater than one, each of the m processing cores capable of supporting a p number of hardware threads, where p is an integer greater than one.

90. (Previously Presented) The method of claim 89, wherein the first abstraction indicates that the first PPM provides an n number of logical processing entities, and wherein the second abstraction indicates that the second PPM provides an s number of logical processing entities, where $s = m$ times p .

91. (Previously Presented) The method of claim 90, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by then logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by at least some of the s logical processing entities of the second PPM.

92. (Previously Presented) The method of claim 91, wherein the resource sharing information of the first abstraction indicates that the n logical processing entities of the first PPM share a cache, and wherein the resource sharing information of the second abstraction indicates that at least some of the s logical processing entities of the second PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.

93. (Previously Presented) The method of claim 26,

wherein the first PPM has a single processing core capable of supporting an n number of hardware threads, where n is an integer greater than one; and

wherein the second PPM has an m number of processing cores, wherein m is an integer greater than one, each of the m processing cores capable of supporting a p number of hardware threads, where p is an integer greater than one.

94. (Previously Presented) The method of claim 93, wherein the first abstraction indicates that the first PPM provides an n number of logical processing entities, and wherein the second abstraction indicates that the second PPM provides an s number of logical processing entities, where $s = m$ times p .
95. (Previously Presented) The method of claim 94, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the n logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by at least some of the s logical processing entities of the second PPM.
96. (Previously Presented) The method of claim 95, wherein the resource sharing information of the first abstraction indicates that the n logical processing entities of the first PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache, and wherein the resource sharing information of the second abstraction indicates that at least some of the s logical processing entities of the second PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.
97. (Currently Amended) The method of claim 26, further comprising:
accessing the first abstraction and/or the second abstraction; and
~~determining, based at least partially upon the operational information in the first abstraction and/or the operational information in the second abstraction, whether to dispatch an execution thread to one of the logical processing entities of the first PPM or one of the logical processing entities of the second PPM to be executed thereby.~~
98. (Previously Presented) The method of claim 97, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the

logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by the logical processing entities of the second PPM.

99. (Previously Presented) The method of claim 98, wherein determining whether to dispatch an execution thread comprises:

determining, based at least partially upon the resource sharing information in the first abstraction and/or the resource sharing information in the second abstraction, whether to dispatch an execution thread to one of the logical processing entities of the first PPM or one of the logical processing entities of the second PPM to be executed thereby.

100. (Previously Presented) The method of claim 99, wherein the resource sharing information in the first abstraction indicates that the logical processing entities of the first PPM are sharing one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.

101. (Previously Presented) The method of claim 100, wherein the resource sharing information in the second abstraction indicates that the logical processing entities of the second PPM are sharing one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.

102. (Previously Presented) The computer readable storage medium of claim 54, wherein the first PPM has an n number of processing cores, where n is an integer greater than one, each of the n processing cores capable of supporting one hardware thread; and wherein the second PPM has a single processing core capable of supporting an m number of hardware threads, where m is an integer greater than one.

103. (Previously Presented) The computer readable storage medium of claim 102, wherein the first abstraction indicates that the first PPM provides an n number of logical processing entities, and wherein the second abstraction indicates that the second PPM provides an m number of logical processing entities.
104. (Previously Presented) The computer readable storage medium of claim 103, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the n logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by the m logical processing entities of the second PPM.
105. (Previously Presented) The computer readable storage medium of claim 104, wherein the resource sharing information of the first abstraction indicates that the n logical processing entities of the first PPM share a cache, and wherein the resource sharing information of the second abstraction indicates that the m logical processing entities of the second PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.
106. (Previously Presented) The computer readable storage medium of claim 54,
wherein the first PPM has an n number of processing cores, where n is an integer greater than one, each of the n processing cores capable of supporting one hardware thread;
and
wherein the second PPM has an m number of processing cores, wherein m is an integer greater than one, each of the m processing cores capable of supporting a p number of hardware threads, where p is an integer greater than one.
107. (Previously Presented) The computer readable storage medium of claim 106, wherein the first abstraction indicates that the first PPM provides an n number of logical processing entities, and wherein the second abstraction indicates that the second PPM provides an s number of logical processing entities, where $s = m$ times p .

108. (Previously Presented) The computer readable storage medium of claim 107, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the n logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by at least some of the s logical processing entities of the second PPM.
109. (Previously Presented) The computer readable storage medium of claim 108, wherein the resource sharing information of the first abstraction indicates that the n logical processing entities of the first PPM share a cache, and wherein the resource sharing information of the second abstraction indicates that at least some of the s logical processing entities of the second PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.
110. (Previously Presented) The computer readable storage medium of claim 54, wherein the first PPM has a single processing core capable of supporting an n number of hardware threads, where n is an integer greater than one; and wherein the second PPM has an m number of processing cores, wherein m is an integer greater than one, each of the m processing cores capable of supporting a p number of hardware threads, where p is an integer greater than one.
111. (Previously Presented) The computer readable storage medium of claim 110, wherein the first abstraction indicates that the first PPM provides an n number of logical processing entities, and wherein the second abstraction indicates that the second PPM provides an s number of logical processing entities, where $s = m$ times p .
112. (Previously Presented) The computer readable storage medium of claim 111, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the n logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information

indicating one or more resources shared by at least some of the s logical processing entities of the second PPM.

113. (Previously Presented) The computer readable storage medium of claim 112, wherein the resource sharing information of the first abstraction indicates that the n logical processing entities of the first PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache, and wherein the resource sharing information of the second abstraction indicates that at least some of the s logical processing entities of the second PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.

114. (Currently Amended) The computer readable storage medium of claim 54, wherein the instructions cause the computer system to further perform the operations of:
accessing the first abstraction and/or the second abstraction; and
~~determining, based at least partially upon the operational information in the first abstraction and/or the operational information in the second abstraction, whether to dispatch an execution thread to one of the logical processing entities of the first PPM or one of the logical processing entities of the second PPM to be executed thereby.~~

115. (Previously Presented) The computer readable storage medium of claim 114, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by the logical processing entities of the second PPM.

116. (Previously Presented) The computer readable storage medium of claim 115, wherein determining whether to dispatch an execution thread comprises:
determining, based at least partially upon the resource sharing information in the first abstraction and/or the resource sharing information in the second abstraction, whether to dispatch an execution thread to one of the logical processing entities of

the first PPM or one of the logical processing entities of the second PPM to be executed thereby.

117. (Previously Presented) The computer readable storage medium of claim 116, wherein the resource sharing information in the first abstraction indicates that the logical processing entities of the first PPM are sharing one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.

118. (Previously Presented) The computer readable storage medium of claim 117, wherein the resource sharing information in the second abstraction indicates that the logical processing entities of the second PPM are sharing one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.

119. (Previously Presented) The computer system of claim 82,
wherein the first PPM has an n number of processing cores, where n is an integer greater than one, each of the n processing cores capable of supporting one hardware thread;
and
wherein the second PPM has a single processing core capable of supporting an m number of hardware threads, where m is an integer greater than one.

120. (Previously Presented) The computer system of claim 119, wherein the first abstraction indicates that the first PPM provides an n number of logical processing entities, and wherein the second abstraction indicates that the second PPM provides an m number of logical processing entities.

121. (Previously Presented) The computer system of claim 120, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the n logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by the m logical processing entities of the second PPM.

122. (Previously Presented) The computer system of claim 121, wherein the resource sharing information of the first abstraction indicates that the n logical processing entities of the first PPM share a cache, and wherein the resource sharing information of the second abstraction indicates that the m logical processing entities of the second PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.
123. (Previously Presented) The computer system of claim 82,
wherein the first PPM has an n number of processing cores, where n is an integer greater than one, each of the n processing cores capable of supporting one hardware thread;
and
wherein the second PPM has an m number of processing cores, wherein m is an integer greater than one, each of the m processing cores capable of supporting a p number of hardware threads, where p is an integer greater than one.
124. (Previously Presented) The computer system of claim 123, wherein the first abstraction indicates that the first PPM provides an n number of logical processing entities, and wherein the second abstraction indicates that the second PPM provides an s number of logical processing entities, where $s = m$ times p .
125. (Previously Presented) The computer system of claim 124, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the n logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by at least some of the s logical processing entities of the second PPM.
126. (Previously Presented) The computer system of claim 125, wherein the resource sharing information of the first abstraction indicates that the n logical processing entities of the first PPM share a cache, and wherein the resource sharing information of the second abstraction indicates that at least some of the s logical processing entities of the second PPM share one or

more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.

127. (Previously Presented) The computer system of claim 82, wherein the first PPM has a single processing core capable of supporting an n number of hardware threads, where n is an integer greater than one; and wherein the second PPM has an m number of processing cores, wherein m is an integer greater than one, each of the m processing cores capable of supporting a p number of hardware threads, where p is an integer greater than one.
128. (Previously Presented) The computer system of claim 127, wherein the first abstraction indicates that the first PPM provides an n number of logical processing entities, and wherein the second abstraction indicates that the second PPM provides an s number of logical processing entities, where $s = m$ times p .
129. (Previously Presented) The computer system of claim 128, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the n logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by at least some of the s logical processing entities of the second PPM.
130. (Previously Presented) The computer system of claim 129, wherein the resource sharing information of the first abstraction indicates that the n logical processing entities of the first PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache, and wherein the resource sharing information of the second abstraction indicates that at least some of the s logical processing entities of the second PPM share one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.

131. (Currently Amended) The computer system of claim 82, wherein the operating system further performs the operations of:
- accessing the first abstraction and/or the second abstraction; and
 - determining, based at least partially upon the operational information in the first abstraction and/or the operational information in the second abstraction, whether to dispatch an execution thread to one of the logical processing entities of the first PPM or one of the logical processing entities of the second PPM to be executed thereby.
132. (Previously Presented) The computer system of claim 131, wherein the operational information of the first abstraction includes resource sharing information indicating one or more resources shared by the logical processing entities of the first PPM, and wherein the operational information of the second abstraction includes resource sharing information indicating one or more resources shared by the logical processing entities of the second PPM.
133. (Previously Presented) The computer system of claim 132, wherein determining whether to dispatch an execution thread comprises:
- determining, based at least partially upon the resource sharing information in the first abstraction and/or the resource sharing information in the second abstraction, whether to dispatch an execution thread to one of the logical processing entities of the first PPM or one of the logical processing entities of the second PPM to be executed thereby.
134. (Previously Presented) The computer system of claim 133, wherein the resource sharing information in the first abstraction indicates that the logical processing entities of the first PPM are sharing one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.
135. (Previously Presented) The computer system of claim 134, wherein the resource sharing information in the second abstraction indicates that the logical processing entities of the second PPM are sharing one or more of the following: a processing core, a data pathway, a translation lookaside buffer, and a cache.